

# Exercises for Day 4 – Solution

## Applied Statistics & Statistical methods for SCIENCE

Anton Rask Lundborg

December 2023

### Exercise 4.1 Multilinear regression and multicollinearity

#### Problem

Multilinear regression refers to the situation where several continuous covariates are used together as explanatory variables in a regression analysis. When doing a multilinear regression you should be aware of the potential pitfalls that may arise if the covariates are *multicollinear*. The purpose of this exercise is to exemplify these pitfalls. This exercise should be done without opening RStudio, but if you want to try the R code yourselves you may find the dataset in the file `wage.txt`.

We consider data taken from *The Current Population Survey (CPS)* made in the US in 1985. The dataset contains observations of the following 6 variables for 532 individuals:

`edu`: length of the person's total education in years.

`sex`: gender of the person's (1=female, 0=male).

`exper`: length of the person's working experience in years.

`wage`: wage in US dollars per hour.

`age`: age of the person in years.

`occup`: profession (1=management, 2=trade, 3=office, 4=service, 5=craft, 6=other).

The following R code fits a multilinear regression of wage on length of education, length of working experience, and age among women working with craftsmanship<sup>1</sup>:

```
> summary(lm(wage ~ edu + exper + age,
             data = subset(wage, (sex == 1) & (occup == 5))))
```

Call:

```
lm(formula = wage ~ edu + exper + age, data = wage_subset)
```

Residuals:

Min	1Q	Median	3Q	Max
-6.5109	-2.9453	-0.6629	2.0672	14.0105

Coefficients:

	Estimate	Std. Error	t value	Pr(> t )
(Intercept)	-15.4931	6.6457	-2.331	0.024 *

---

<sup>1</sup>Model validation would reveal that it is better to model the logarithm of the wage but in order not to give the impression that responses should always be log-transformed (which is true!), and also to keep the interpretations of the parameters estimates as simple as possible, we will not transform the response variable. This is ok since the emphasis of this exercise is *multicollinearity* and not *model validity*.

```

edu          0.7059    0.8524    0.828    0.412
exper       -0.6247    0.8723   -0.716    0.477
age          0.6775    0.7964    0.851    0.399
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

```

```

Residual standard error: 4.593 on 48 degrees of freedom
Multiple R-squared:  0.3097,    Adjusted R-squared:  0.2666
F-statistic: 7.179 on 3 and 48 DF,  p-value: 0.0004461

```

From the model summary we see that neither of the 3 explanatory variables are close to significance. However, the multilinear regression still explains 30.97% (i.e. the  $R^2$ ) of the variation in the wages, and taken together the 3 explanatory variables are highly significant ( $p = 0.0004461$ ).

- Compare the **summary**-output to the statements made above and confirm that the reporting of hypothesis tests and  $R^2$  is correct.
- Is there an interpretation of the sign of the estimates for the 3 slopes? E.g. do craftswomen earn more if they have more working experience? Or is it impossible to make such an interpretation in this case?
- An automated backward model reduction would proceed by removing **exper** being the least significant variable. However, what are the arguments for removing **age** instead?

The fit of the multilinear regression after removal of **age** is given on the next page. Please consider the following questions:

- What has happened to the p-values for **edu** and **exper**?
- What has happened to the sign of the slope of **exper**? Do you think that the positive sign makes more sense? Why/why not?

```

> summary(lm(wage ~ edu + exper,
             data = subset(wage, (sex == 1) & (occup == 5))))

```

```

Call:
lm(formula = wage ~ edu + exper, data = subset(wage, (sex ==
1) & (occup == 5)))

```

```

Residuals:
    Min       1Q   Median       3Q      Max
-5.9828 -3.0854 -0.6495  1.7550 14.1748

```

```

Coefficients:
            Estimate Std. Error t value Pr(>|t|)
(Intercept) -11.85350    5.07118  -2.337  0.0235 *
edu           1.38007    0.31307   4.408 5.68e-05 ***
exper         0.11552    0.06237   1.852  0.0700 .
---

```

```

Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

```

```

Residual standard error: 4.581 on 49 degrees of freedom
Multiple R-squared:  0.2993,    Adjusted R-squared:  0.2707
F-statistic: 10.47 on 2 and 49 DF,  p-value: 0.0001641

```

The following output from R shows that **edu** and **exper** may be considered uncorrelated in the subpopulation of craftswomen. Does this have any implication for the interpretation of the slope estimates on **edu** and **exper** given above? Why/why not? And what if **edu** and **exper** actually are negatively correlated, i.e. if working experience in general is shorter for craftswomen with a longer education?

```
> with(subset(wage, (sex == 1) & (occup == 5)), cor.test(educ, exper))
```

Pearson's product-moment correlation

```
data:  educ and exper
t = -1.0381, df = 50, p-value = 0.3042
alternative hypothesis: true correlation is not equal to 0
95 percent confidence interval:
 -0.4022055  0.1329212
sample estimates:
      cor
-0.1452478
```

Multicollinearity means that some of the covariates explains the same property in the experimental units e.g. if you have a long education as well as long working experience, then you will most likely also have a comparably high age. It stands to reason that we will only need two of the three variables `educ`, `exper`, `age` in order to characterize these properties of a person. To decide which two of these variables provides the “correct” explanation can not be done based on statistics, but relies on the interpretation of the variables. When there is multicollinearity among the explanatory variables, the  $p$ -values may change from non-significant to highly significant and the estimates may change sign after model reduction. That there is multicollinearity in the present dataset may be seen from the following analysis<sup>2</sup>

```
> summary(lm(age ~ educ + exper, data = wage))
```

```
Call: lm(formula = age ~ educ + exper, data = wage)
```

```
Residuals:
```

```
      Min       1Q   Median       3Q      Max
-3.8507 -0.3801 -0.0122  0.4081  2.1230
```

```
Coefficients:
```

```
              Estimate Std. Error t value Pr(>|t|)
(Intercept)  6.09160     0.19182   31.76  <2e-16 ***
educ         0.98494     0.01281   76.91  <2e-16 ***
exper       1.05558     0.00271  389.51  <2e-16 ***
```

```
---
```

```
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

```
Residual standard error: 0.7235 on 529 degrees of freedom
```

```
Multiple R-squared:  0.9966,    Adjusted R-squared:  0.9966
```

```
F-statistic: 7.793e+04 on 2 and 529 DF,  p-value: < 2.2e-16
```

Please do the following:

- Comment on the  $R^2$ -value as well as the significance tests.
- What is the interpretation of the estimate of the intercept?
- What is the interpretation of the null hypothesis that the slopes on `educ` and `exper` both equal 1?
- What is the interpretation of the error term, and the  $RMSE = 0.7235$ ?

## Solution

We first load the data and reproduce the summary.

```
wage <- read.delim("wage.txt")
summary(lm(wage ~ educ + exper + age,
```

---

<sup>2</sup>A few more suggestions for the identification of multicollinearity may be found in the solution to the exercise.

```

data = subset(wage, (sex == 1) & (occup == 5)))

##
## Call:
## lm(formula = wage ~ edu + exper + age, data = subset(wage, (sex ==
## 1) & (occup == 5)))
##
## Residuals:
##      Min       1Q   Median       3Q      Max
## -6.5109 -2.9453 -0.6629  2.0672 14.0105
##
## Coefficients:
##              Estimate Std. Error t value Pr(>|t|)
## (Intercept) -15.4931     6.6457  -2.331  0.024 *
##      edu         0.7059     0.8524   0.828  0.412
##     exper       -0.6247     0.8723  -0.716  0.477
##      age         0.6775     0.7964   0.851  0.399
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Residual standard error: 4.593 on 48 degrees of freedom
## Multiple R-squared:  0.3097, Adjusted R-squared:  0.2666
## F-statistic: 7.179 on 3 and 48 DF,  p-value: 0.0004461

```

We see that despite the fact that none of the covariates are significant, the  $R^2$  is around 0.31 and that taken together the three variables are highly significant with a  $p$ -value around 0.0004. This model is very difficult to interpret though as any increase in `exper` or `edu` necessarily also result in an increase in `age` and for most people we would expect `exper + edu` to be equal to around `age` (minus some fixed intercept). While the  $p$ -value of `exper` is the largest and thus would be automatically omitted by an algorithm, it seems more natural to remove `age` as we can at least think of education and work experience as distinct quantities. We fit this model:

```

summary(lm(wage ~ edu + exper, data = subset(wage, (sex == 1) & (occup == 5))))

##
## Call:
## lm(formula = wage ~ edu + exper, data = subset(wage, (sex ==
## 1) & (occup == 5)))
##
## Residuals:
##      Min       1Q   Median       3Q      Max
## -5.9828 -3.0854 -0.6495  1.7550 14.1748
##
## Coefficients:
##              Estimate Std. Error t value Pr(>|t|)
## (Intercept) -11.85350     5.07118  -2.337  0.0235 *
##      edu         1.38007     0.31307   4.408 5.68e-05 ***
##     exper         0.11552     0.06237   1.852  0.0700 .
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Residual standard error: 4.581 on 49 degrees of freedom
## Multiple R-squared:  0.2993, Adjusted R-squared:  0.2707
## F-statistic: 10.47 on 2 and 49 DF,  p-value: 0.0001641

```

We see that the  $p$ -values of both `edu` and `exper` have decreased significantly. In addition, the slope of `exper`

has become positive which is closer to what one would expect; increases in work experience usually lead to higher wages. We can check that the remaining two covariates are uncorrelated:

```
with(subset(wage, (sex == 1) & (occup == 5)), cor.test(educ, exper))

##
## Pearson's product-moment correlation
##
## data:  educ and exper
## t = -1.0381, df = 50, p-value = 0.3042
## alternative hypothesis: true correlation is not equal to 0
## 95 percent confidence interval:
## -0.4022055  0.1329212
## sample estimates:
##          cor
## -0.1452478
```

The fact that we can accept the hypothesis that these covariates are independent means that we can freely interpret the slopes independently of each other. If the correlation was negative, then we would have to keep in mind that an increase in, say, `exper` would usually result in a decrease of `educ` thus we could not view each slope estimate in isolation. We can show that the original three variables were highly collinear by performing the following analysis:

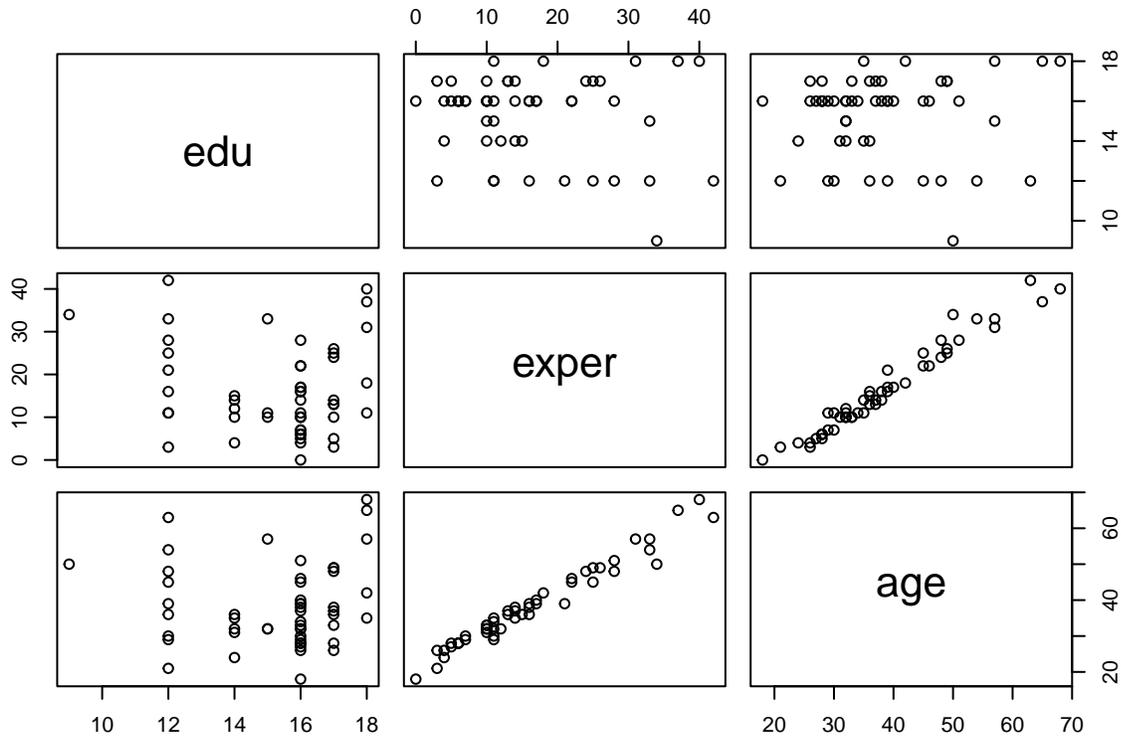
```
summary(lm(age ~ educ + exper, data = wage))

##
## Call:
## lm(formula = age ~ educ + exper, data = wage)
##
## Residuals:
##      Min       1Q   Median       3Q      Max
## -3.8507 -0.3801 -0.0122  0.4081  2.1230
##
## Coefficients:
##              Estimate Std. Error t value Pr(>|t|)
## (Intercept)  6.09160    0.19182   31.76  <2e-16 ***
## educ         0.98494    0.01281   76.91  <2e-16 ***
## exper        1.05558    0.00271  389.51  <2e-16 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Residual standard error: 0.7235 on 529 degrees of freedom
## Multiple R-squared:  0.9966, Adjusted R-squared:  0.9966
## F-statistic: 7.793e+04 on 2 and 529 DF,  p-value: < 2.2e-16
```

We see that `educ` and `exper` almost perfectly predict `age` ( $R^2 = 0.9966$ ) and that both variables are required for the prediction (as the  $p$ -values are tiny). The intercept estimate tells us that a good estimate of the number of years without partaking in work or education is around 6. The hypothesis that both slopes are equal to one correspond to the hypothesis that age is equal to a fixed number plus the number of years of education plus the number of years of work experience, that is, that working and studying do not happen simultaneously. The error term tells us about the deviation of each individual from the population average and the RMSE of 0.7235 is an estimate of the standard deviation of the error term. If the errors were normally distributed, we would expect 95% of ages to lie within  $\pm 1.96 \cdot 0.7235 = 1.418$  of the population average.

An alternative way of spotting correlation between variables is to plot pairs of explanatory variables and computing their correlations:

```
pairs(wage[(wage$sex == 1) & (wage$occup == 5)], c("educ", "exper", "age"))
```



```
cor(wage[(wage$sex == 1) & (wage$occup == 5), c("edu", "exper", "age")])
```

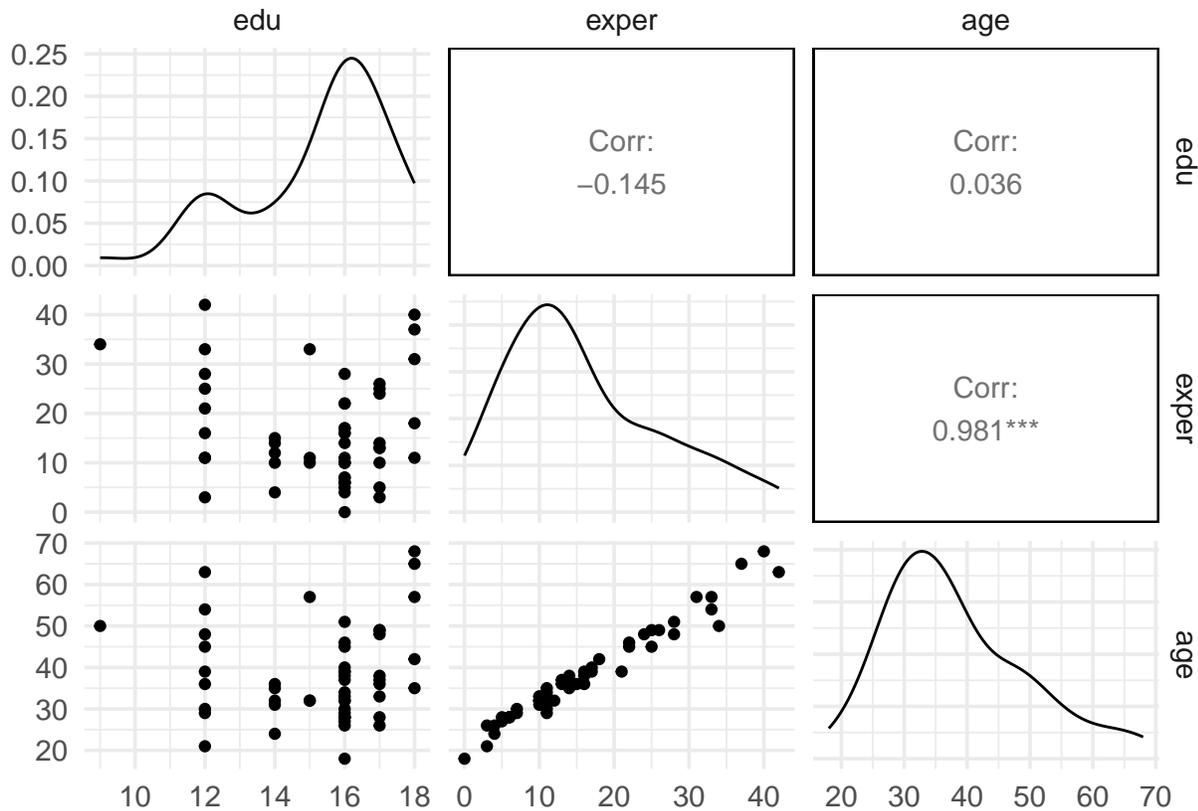
```
##           edu      exper      age
## edu      1.0000000 -0.1452478  0.03645816
## exper -0.14524783  1.0000000  0.98089765
## age      0.03645816  0.9808977  1.00000000
```

Here, it is already clear that `age` and `exper` are highly correlated. There is a fancy `ggplot` version of this plot in the `GGally` package:

```
library(GGally)

## Registered S3 method overwritten by 'GGally':
##   method from
##   +.gg      ggplot2

ggpairs(wage[(wage$sex == 1) & (wage$occup == 5), c("edu", "exper", "age")])
```



It can be harder to spot correlation between 3 or more variables from these plots. One way to catch this is to perform a Principal Components Analysis (PCA):

```
summary(
  prcomp(wage[(wage$sex == 1) & (wage$occup == 5)], c("edu", "exper", "age")),
  scale = TRUE
)

## Importance of components:
##              PC1    PC2    PC3
## Standard deviation  1.4096 1.0053 0.05003
## Proportion of Variance 0.6623 0.3368 0.00083
## Cumulative Proportion 0.6623 0.9992 1.00000
```

The cumulative proportion tell us how much variance in the data is explained with 1, 2 or 3 variables, respectively. We see that more than 99.9% of variance is explained by just using two variables (which indicates multicollinearity).

## Exercise 4.2 ANCOVA, statistical modeling, and more

### Problem

In this exercise we investigate the world records for outdoor running distances. The records were taken from the website <http://www.iaaf.org> of the International Association of Athletics Federation on May 7, 2011. We want to examine the dependence of the record (**time**) on the **distance**, and to examine the difference between men and women. The purpose of this exercise is to give a non-trivial example of the choices needed in making simple statistical models with sensible interpretations.

The following items guide you through such an analysis step-by-step:

- Read the dataset available in the text file `WR2011.txt` into R (in a data frame called `wr`), and have a look at the variables:
  - Please note that the distances are more or less doubled between consecutive running disciplines. Thus, the running distances are almost equidistant on a logarithmic scale.
  - The variable `DOB` contains the data-of-birth of the record holder. The variables `Place` and `Date` contain the place and date of the record. These variables will not be used in this exercise.
  - The variable `bend` is not in the original data, and will be used later. This variable quantifies how many times longer than 1500 meters the running distance in question is, and it is set to 1 if the distance is shorter than 1500 meters.
  - Make sure that the variables `time`, `distance` and `bend` are numerical, and that `sex` is a categorical factor.
- Produce a plot of `time` against `distance` using the code:

```
library(ggplot2)
ggplot(wr) + geom_point(aes(x = distance, y = time, col = sex))
```

This plot corresponds to the relationship:

$$\text{time} = \alpha + \beta \cdot \text{distance}.$$

The parameter  $\beta$  describes the running velocity. Thus, in this model the running velocity is the same no matter the distance. Is this realistic?

- Create a plot of  $\log(\text{time})$  against  $\log(\text{distance})$ . This plot corresponds to the relationship:

$$\log(\text{time}) = \alpha + \beta \cdot \log(\text{distance}).$$

Taking the exponential function on both sides we find:

$$\text{time} = \exp(\alpha) \cdot \log(\text{distance})^\beta.$$

It is not obvious that this is a good model. But what do you think looking at the plot?

- Perform the following linear regression:

```
m1 <- lm(log(time) ~ sex + log(distance) + sex:log(distance), data = wr)
```

In this model both the  $\alpha$  and the  $\beta$  parameter depend on the gender:

$$\log(\text{time}_i) = \alpha(\text{sex}_i) + \beta(\text{sex}_i) \cdot \log(\text{distance}_i) + \text{error}_i.$$

Such a model is called an ANCOVA (ANalysis of COVariance). The ANCOVA will allow us to compare the records for men and women.

- Is this ANCOVA model valid? Look in particular at the *residual plot* (called “Residuals vs Fitted” if you use `plot(m1)`). It seems there is a bend at observations number 6 and 22. One way to identify these observation numbers is to use the `identify()` function. To use this the residual plot should fill the entire graphics window, so we will make it again. Try the R code<sup>3</sup>

```
par(mfrow=c(1, 1))
plot(predict(m1), residuals(m1))
identify(predict(m1), residuals(m1))
```

and use the mouse to click on the points where you think the bend is positioned. After you are done, finish the identifier as signified in the graphics window (in Windows you should press the *Esc*-key).

Remark: If `identity()` does not work inside RStudio, then a solution might be to open a separate graphical device using the function `x11()` before making the plot.

---

<sup>3</sup>The line `par(mfrow=c(1, 1))` is only necessary if you did `par(mfrow=c(2, 2))` before.

- Check that observations number 6 and 22 correspond to the 1500-meter distance for men and women, respectively. Thus, there appears to be a difference between short running distances (less than 1500 meters) and long running distances (more than 1500 meters).
- In order to allow the regression lines to bend at 1500 meters we include the variable `bend` in the ANCOVA model:

$$\log(\text{time}_i) = \alpha(\text{sex}_i) + \beta(\text{sex}_i) \cdot \log(\text{distance}_i) + \gamma(\text{sex}_i) \cdot \log(\text{bend}_i) + \text{error}_i.$$

Fit this model.

- Is the extended ANCOVA model valid? (Hint: No.)
- Refit the extended ANCOVA without using observations number 1, 11, 12, 14, 15, 17, 27, 28, 30, and 31. This may be done using the option

```
data = wr[-c(1, 11, 12, 14, 15, 17, 27, 28, 30, 31),]
```

in the call to `lm()`. Does this improve the model validity?

- Which distances do the removed observations represent? Do you think it is fair to remove the world records for these running distances from the present analysis? Why/why not?

Hint: This may require some knowledge about athletics.

In any case, for the remainder of this exercise you should remove these observations from the analysis.

- Use the function `step()` to do model selection based on the Akaike Information Criterion.
- The AIC-based selection should result in the model, where the interactions between `sex` and `log(distance)`, `log(bend)` are removed:

$$\log(\text{time}_i) = \alpha(\text{sex}_i) + \beta \cdot \log(\text{distance}_i) + \gamma \cdot \log(\text{bend}_i) + \text{error}_i.$$

Thus, the difference between men and women is only via the alpha parameter. Convince yourself that the number

$$\exp(\alpha(\text{woman}) - \alpha(\text{man}))$$

quantifies how much slower the women run than the men.

- Give a 95% confidence interval for the relative running speed of women compared to men.

Hint: The contrast  $\alpha(\text{woman}) - \alpha(\text{man})$  is called `sexwoman` in the parametrization used by R.

(Reference: Based on exercise 8.2 from Anders Tolver & Helle Sørensen: *Lecture notes for Applied Statistics*.)

## Solution

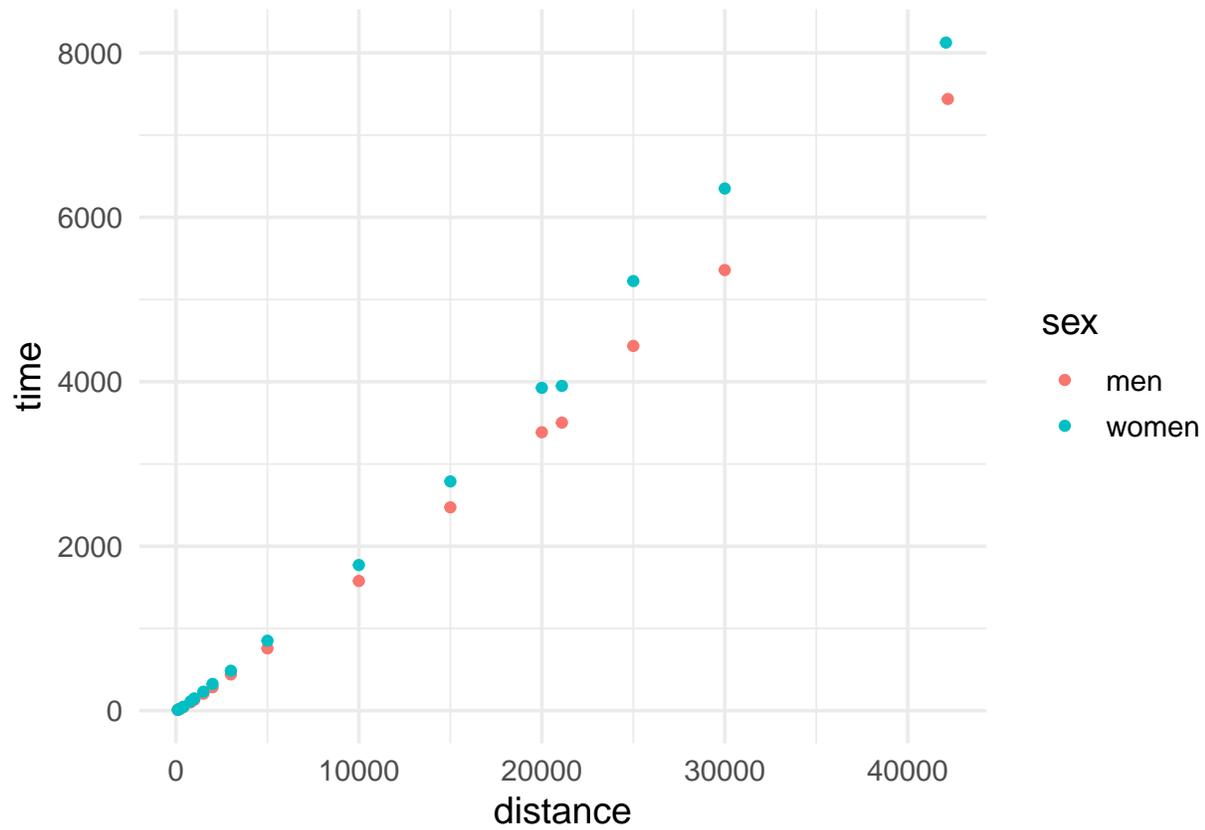
We load the data and check the datatypes:

```
wr <- read.delim("WR2011.txt")
str(wr)

## 'data.frame':   32 obs. of  9 variables:
## $ time       : num  9.58 19.19 43.18 101.01 131.96 ...
## $ distance   : num  100 200 400 800 1000 1500 2000 3000 5000 10000 ...
## $ bend       : num  1 1 1 1 1 ...
## $ sex        : chr  "men" "men" "men" "men" ...
## $ name       : chr  "Usain.Bolt" "Usain.Bolt" "Michael.Johnson" "David.Lekuta.Rudisha" ...
## $ DOB        : chr  "21-08-1986" "21-08-1986" "13-09-1967" "17-12-1988" ...
## $ nationality: chr  "JAM" "JAM" "USA" "KEN" ...
## $ place      : chr  "Berlin" "Berlin" "Sevilla" "Rieti" ...
## $ date       : chr  "16-08-2009" "20-08-2009" "26-08-1999" "29-08-2010" ...
```

Everything looks correct and we plot distance against time:

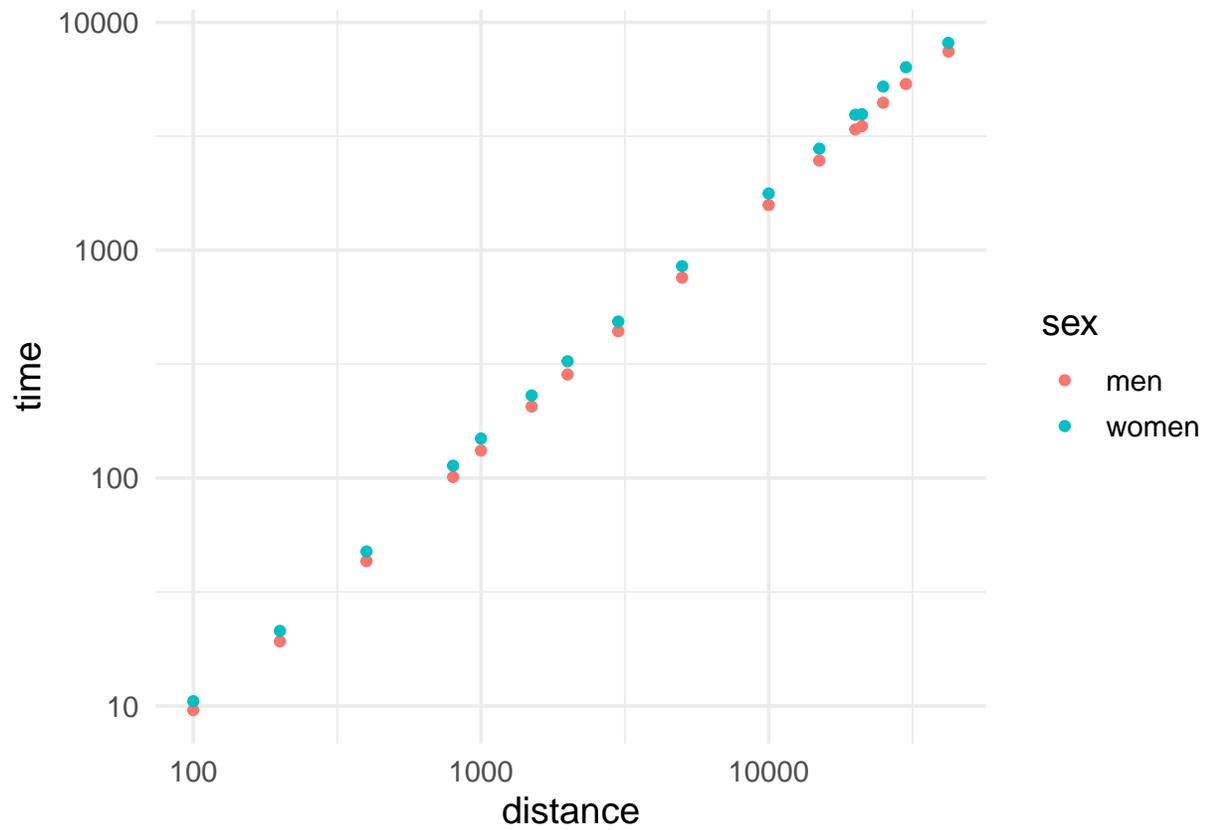
```
library(ggplot2)
ggplot(wr) +
  geom_point(aes(x = distance, y = time, col = sex))
```



This relationship would correspond to a velocity that is constant for all distances. This is clearly not realistic!

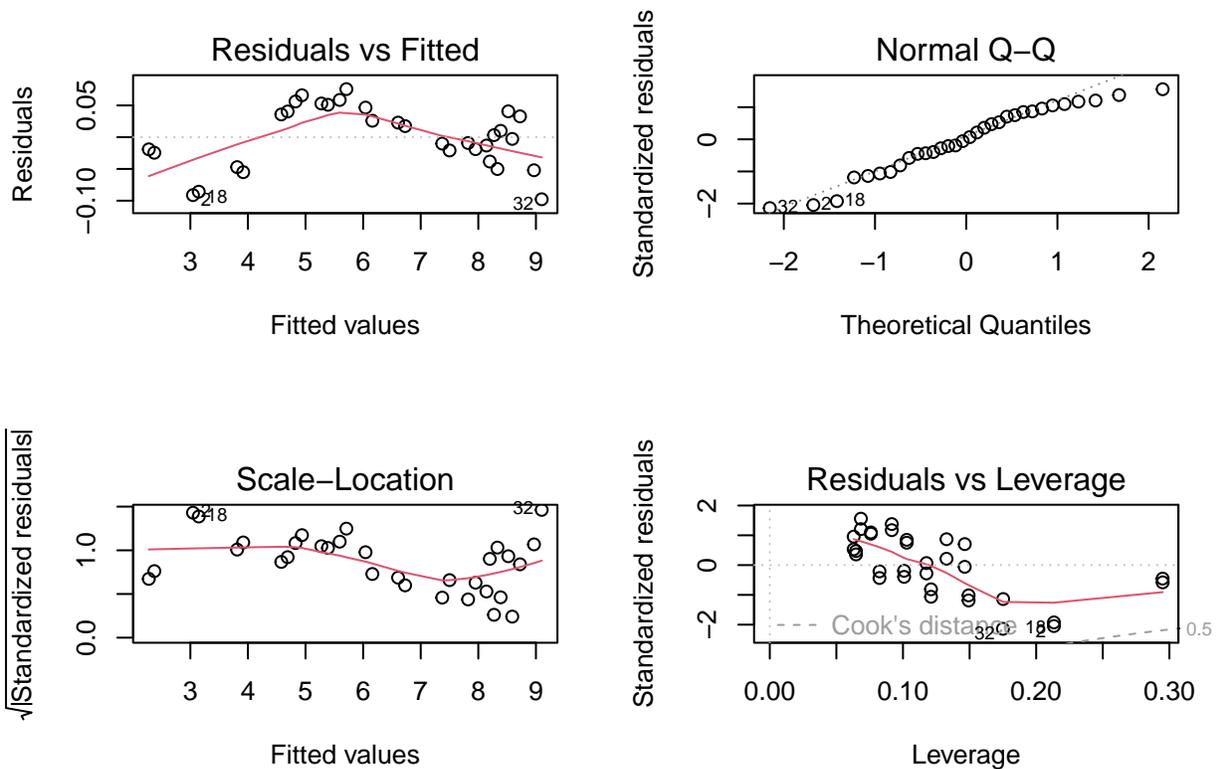
We can instead plot the logarithm of time against the logarithm of distance:

```
ggplot(wr) + geom_point(aes(x = distance, y = time, col = sex)) +
  scale_x_log10() + scale_y_log10()
```



From looking at the plot, it does seem like a decent model; the points look like they are on a line. We can create the model and validate it:

```
m1 <- lm(log(time) ~ sex + log(distance) + sex:log(distance), data = wr)
par(mfrow = c(2, 2))
plot(m1)
```



```
par(mfrow = c(1, 1))
```

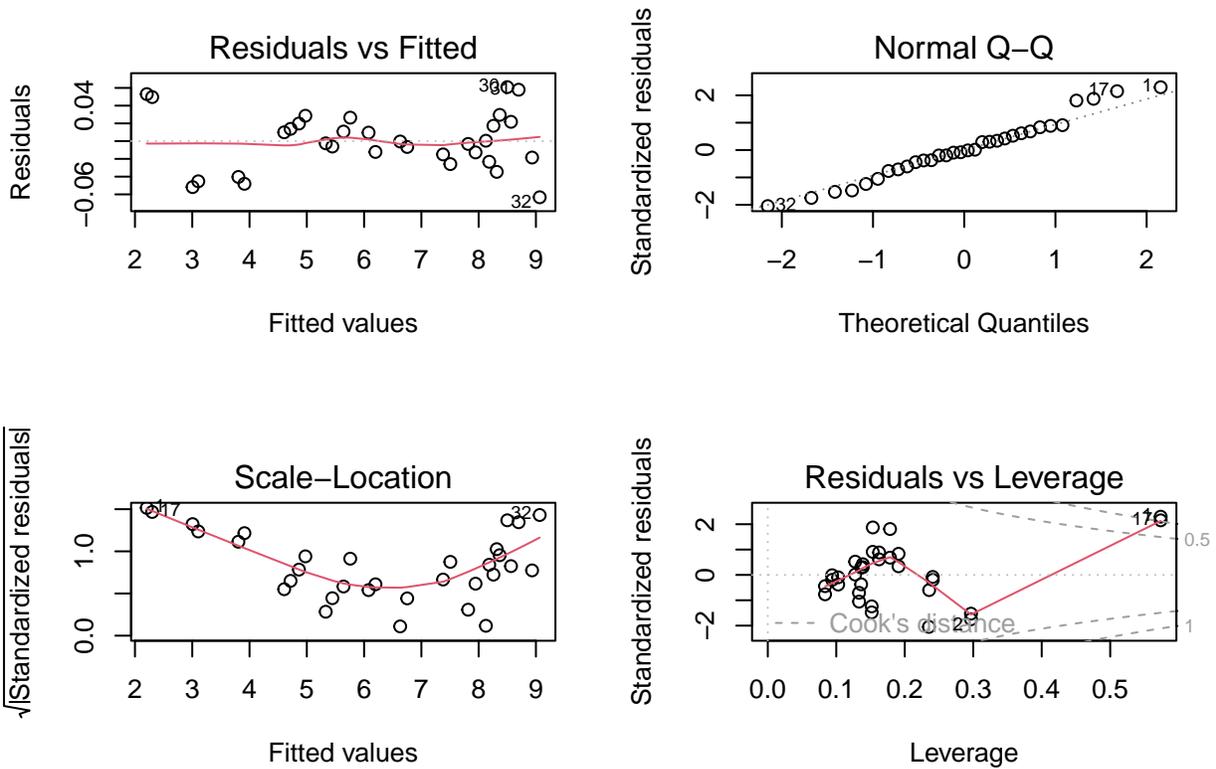
The residuals are not mean zero and there is a clear bend in the values, i.e. the values increase until a certain point after which point they decrease. The points where the bend occur are observations 6 and 22:

```
wr[c(6, 22), ]
```

```
##      time distance bend  sex      name      DOB nationality  place
## 6  206.00    1500    1  men Hicham.El.Guerrouj 14-09-1974      MAR    Roma
## 22 230.46    1500    1 women      Yunxia.Qu 25-12-1972      CHN Beijing
##      date
## 6  14-07-1998
## 22 11-09-1993
```

These are exactly the 1500m distances, so it seems like we have to model observations below 1500m and above 1500m differently. We have encoded this in the `bend` variable that we can include in the model:

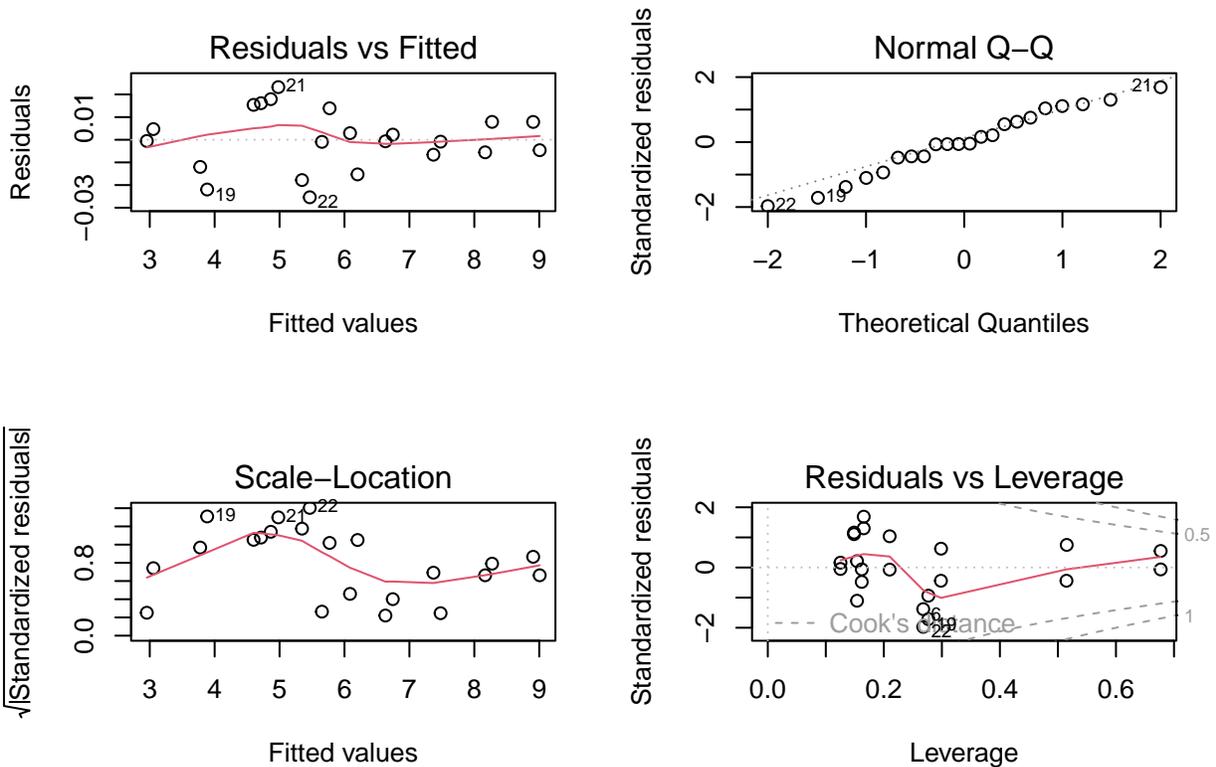
```
m2 <- lm(
  log(time) ~ sex + log(distance) + log(bend) +
  sex:log(distance) + sex:log(bend), data = wr
)
par(mfrow = c(2, 2))
plot(m2)
```



```
par(mfrow = c(1, 1))
```

The plot looks much better but there are still some issues; the variance seems to change for different fitted values and the QQ-plot also does not look that great. We refit the model without some observations:

```
m3 <- lm(
  log(time) ~ sex + log(distance) + log(bend) +
  sex:log(distance) + sex:log(bend),
  data = wr[-c(1, 11, 12, 14, 15, 17, 27, 28, 30, 31), ]
)
par(mfrow = c(2, 2))
plot(m3)
```



```
par(mfrow = c(1, 1))
```

This looks better! We have gotten rid of the worst of the problems with the variance and in the QQ-plot. Let us investigate the distances that we removed:

```
wr[c(1, 11, 12, 14, 15, 17, 27, 28, 30, 31), ]
```

##	time	distance	bend	sex	name	DOB
## 1	9.58	100	1.00000	men	Usain.Bolt	21-08-1986
## 11	2473.00	15000	10.00000	men	Leonard.Patrick.Komon	10-01-1988
## 12	3386.00	20000	13.33333	men	Haile.Gebrselassie	18-04-1973
## 14	4435.80	25000	16.66667	men	Toshihiko.Seko	15-07-1956
## 15	5358.80	30000	20.00000	men	Toshihiko.Seko	15-07-1956
## 17	10.49	100	1.00000	women	Florence.Griffith-Joyner	21-12-1959
## 27	2788.00	15000	10.00000	women	Tirunesh.Dibaba	01-06-1985
## 28	3926.00	20000	13.33333	women	Tegla.Loroupe	09-05-1973
## 30	5225.00	25000	16.66667	women	Tegla.Loroupe	09-05-1973
## 31	6350.00	30000	20.00000	women	Tegla.Loroupe	09-05-1973

##	nationality	place	date
## 1	JAM	Berlin	16-08-2009
## 11	KEN	Nijmegen	21-11-2010
## 12	ETH	Ostrava	27-06-2007
## 14	JPN	Christchurch	22-03-1981
## 15	JPN	Christchurch	22-03-1981
## 17	USA	Indianapolis	16-07-1988
## 27	ETH	Nijmegen	15-11-2009
## 28	KEN	Borgholzhausen	03-09-2000
## 30	KEN	Mengerskirchen	21-09-2002
## 31	KEN	Warstein	06-06-2003

We see that the excluded observations correspond to very short distances (100m) or very long distances

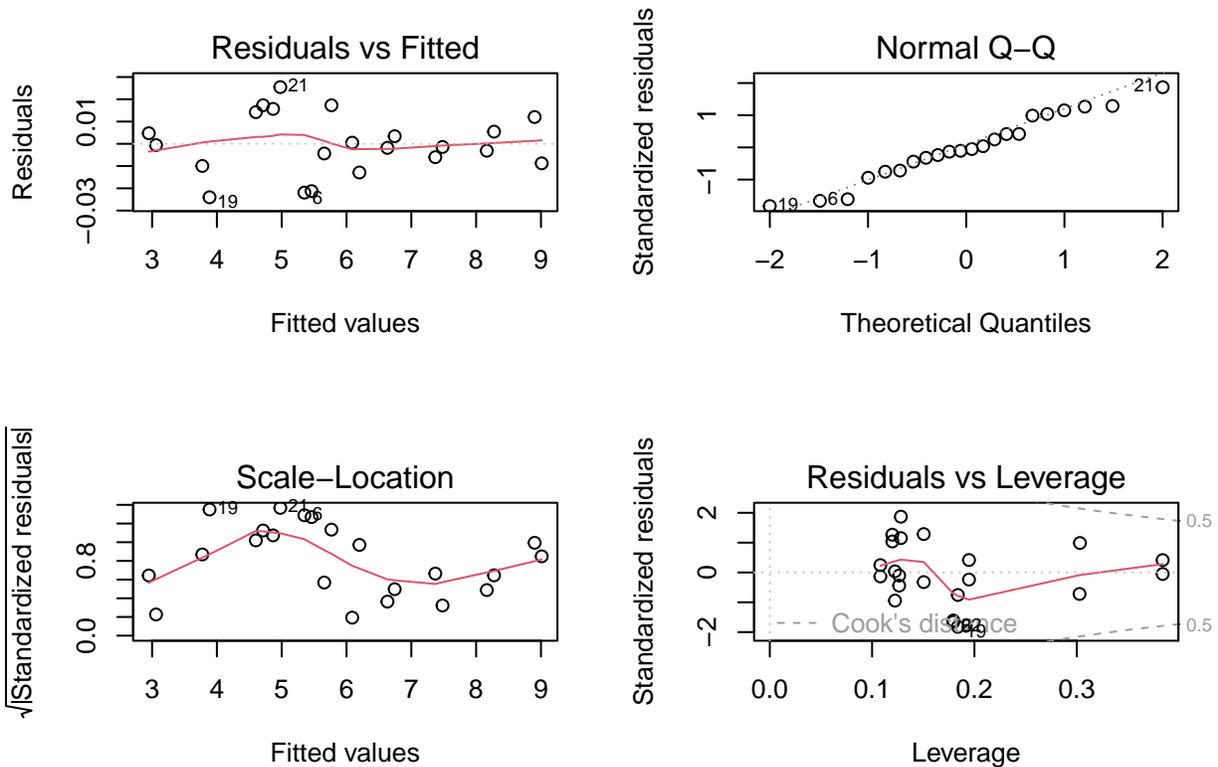
(greater than 15km). It is perhaps reasonable that these are treated differently. We do model selection using the `step()` function:

```
step(m3)

## Start:  AIC=-179.6
## log(time) ~ sex + log(distance) + log(bend) + sex:log(distance) +
##   sex:log(bend)
##
##           Df Sum of Sq      RSS      AIC
## - sex:log(distance)  1 0.00015434 0.0037857 -180.69
## - sex:log(bend)      1 0.00019360 0.0038249 -180.46
## <none>                                0.0036313 -179.60
##
## Step:  AIC=-180.69
## log(time) ~ sex + log(distance) + log(bend) + sex:log(bend)
##
##           Df Sum of Sq      RSS      AIC
## - sex:log(bend)  1      0.000  0.0038 -182.446
## <none>                                0.0038 -180.687
## - log(distance)  1   10.016 10.0202   -9.302
##
## Step:  AIC=-182.45
## log(time) ~ sex + log(distance) + log(bend)
##
##           Df Sum of Sq      RSS      AIC
## <none>                                0.0038 -182.446
## - log(bend)      1    0.0605  0.0643 -122.380
## - sex            1    0.0684  0.0722 -119.813
## - log(distance)  1   10.0164 10.0202  -11.302
##
## Call:
## lm(formula = log(time) ~ sex + log(distance) + log(bend), data = wr[-c(1,
##   11, 12, 14, 15, 17, 27, 28, 30, 31), ])
##
## Coefficients:
##   (Intercept)      sexwomen  log(distance)      log(bend)
##      -3.3618         0.1115         1.1912        -0.1265
```

We investigate the final selected model (which only includes main effects of `sex`, `log(bend)` and `log(distance)`):

```
m4 <- lm(log(time) ~ sex + log(distance) + log(bend),
  data = wr[-c(1, 11, 12, 14, 15, 17, 27, 28, 30, 31), ]
)
par(mfrow = c(2, 2))
plot(m4)
```



```
par(mfrow = c(1, 1))
```

The diagnostics still look okay! We can investigate whether all the remaining variables are important using the `drop1` function:

```
drop1(m4, test = "F")

## Single term deletions
##
## Model:
## log(time) ~ sex + log(distance) + log(bend)
##           Df Sum of Sq   RSS   AIC  F value    Pr(>F)
## <none>                0.0038 -182.446
## sex                1   0.0684  0.0722 -119.813   321.76 6.248e-13 ***
## log(distance)      1  10.0164 10.0202  -11.302 47107.84 < 2.2e-16 ***
## log(bend)          1   0.0605  0.0643  -122.380   284.35 1.791e-12 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

The  $p$ -values are tiny for all predictors and thus we conclude that all of them are important. We can give confidence intervals for the coefficients:

```
cbind(estimate = coef(m4), confint(m4))

##           estimate      2.5 %      97.5 %
## (Intercept) -3.3617597 -3.43932831 -3.2841911
## sexwomen      0.1115313  0.09846844  0.1245941
## log(distance)  1.1912083  1.17967773  1.2027389
## log(bend)     -0.1265301 -0.14229459 -0.1107655
```

To get a confidence interval for the difference in the sexes, we compute the exponential of the effect of `sex` and transform the confidence interval:

```
exp(cbind(women_vs_men = coef(m4), confint(m4))[2, ])
```

```
## women_vs_men      2.5 %      97.5 %
##      1.117989      1.103480      1.132689
```

We see that women run about 11.8% slower than men. We could also have used the `emmeans`-package:

```
library(emmeans)
pairs(emmeans(m4, ~sex), type = "response", reverse = TRUE)
```

```
## contrast      ratio      SE df null t.ratio p.value
## women / men  1.12 0.00695 18   1  17.938 <.0001
##
## Tests are performed on the log scale
```

## Exercise 4.3 Incomplete block experiment

### Problem

The following experiment was carried out by H. Wolffhechel, KVL, in 1986. The purpose of the experiment was to compare 12 sphagnum lots with respect to water and air content. Each lot was applied to four pots with small cucumber plants. The pots were placed in one of six watering troughs, each containing eight pots. The experimental design and the volume (water and air content), in percent, is given in the following table (dataset available in file `sphagnum.txt`) for each pot.

Sphagnum lot	Watering trough						Mean
	1	2	3	4	5	6	
1	37.0		44.6		42.5	47.1	42.80
2			49.0	50.5	51.0	44.8	48.83
3		34.6	42.7	41.8	37.8		39.23
4	45.3	42.7	47.7	42.8			44.63
5	32.1	38.5		32.0		31.6	33.55
6	34.3	33.3		34.0	22.6		31.05
7	32.3		28.1	28.1	32.3		30.20
8	38.9	36.5		39.7		34.8	37.48
9	33.9		31.4	32.1		23.0	30.10
10		39.7	41.8		43.5	33.8	39.70
11		41.1	38.1		31.1	37.9	37.05
12	35.9	7.5			36.2	25.5	26.28

Construct and validate the additive model for `volume`, i.e. the model including only the main effects of the two explanatory variables `lot` and `trough`:

- Remember that the variables `lot` and `trough` should be used as factors. You can achieve this by using the `factor()` function in the call to `lm()`. However, in this exercise it is recommended that you change the type of the variables in the data frame at the beginning of your R code:

```
sphagnum      <- read.delim("sphagnum.txt")
sphagnum$lot  <- factor(sphagnum$lot)
sphagnum$trough <- factor(sphagnum$trough)
```

- You probably want to remove the “extreme” observation `volume = 7.5` for `(lot, trough) = (12, 2)`. Can you use the validation plots in R to identify the number of this observation?

Find the estimated marginal means of `volume` for the 12 different sphagnum lots using the additive model. Why are these em-means different from the raw means listed in the above table? Do you prefer the raw means or the em-means? Why?

Remark: The `emmeans`-package may be used to compute and compare the em-means. The statistical

computations done in the `emmeans`-package are based on standard errors extracted from the model objects. Suppose e.g. that your model is available in an `lm`-object called `m2`, and try the following R code (and think about what the code does):

```
# load library
library(emmeans)

# Create and plot em-means
emmeans(m2, ~ lot)
plot(emmeans(m2, ~ lot))

# Tukey grouping of em-means
# Note: the p-values are adjusted for multiple testing,
#       but the confidence intervals are not adjusted!
# Note: Also needs multcomp-package to be installed
#       (but not necessarily loaded!)
multcomp::cld(emmeans(m2, ~ lot))

# Remark: The author of the emmeans-package, Russell Lenth,
#         does not like the "compact letter display".
#         Earlier there was a CLD() in the emmeans-package, but
#         this functionality has been removed from the package!
#         Luckily, you may use multcomp::cld() instead!

# As a replacement for the CLD() functionality Russell
# proposes the following plot. Some may find this
# display to be too busy. But what do you think?
pwpp(emmeans(m2, ~ lot))

# An alternative to pwpp() is to find and plot
# simultaneous confidence intervals.
# Note, however, that replacing hypothesis tests by
# looking for overlap between confidence intervals
# may be misleading.
confint(emmeans(m2, ~ lot), adjust="tukey")
plot(confint(emmeans(m2, ~ lot), adjust="tukey"))
```

Remark: An alternative is to use the `multcomp` package. However, the syntax for that package can be much more difficult to learn. It is therefore recommended to use the `emmeans` package.

## Solution

We load the data

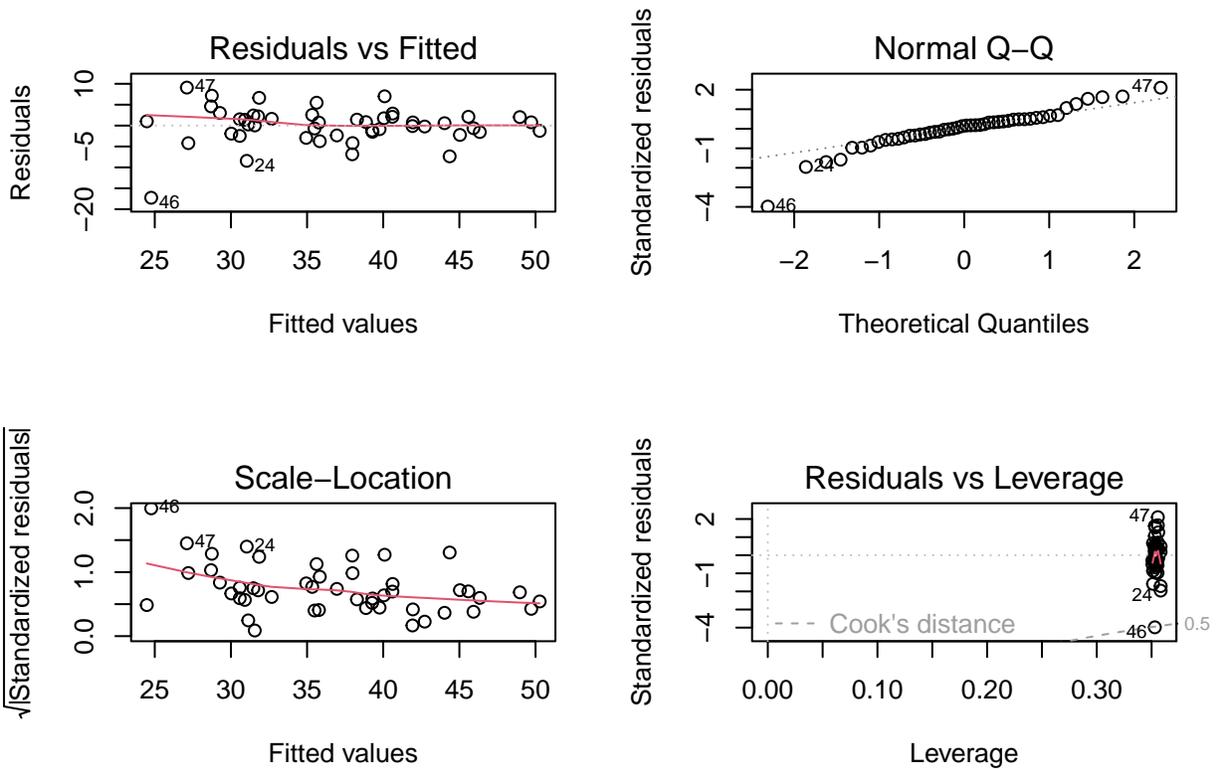
```
sphagnum      <- read.delim("sphagnum.txt")
sphagnum$lot  <- factor(sphagnum$lot)
sphagnum$trough <- factor(sphagnum$trough)
str(sphagnum)

## 'data.frame':   48 obs. of  3 variables:
## $ lot      : Factor w/ 12 levels "1","2","3","4",...: 1 1 1 1 2 2 2 2 3 3 ...
## $ trough: Factor w/ 6 levels "1","2","3","4",...: 1 3 5 6 3 4 5 6 2 3 ...
## $ volume: num  37 44.6 42.5 47.1 49 50.5 51 44.8 34.6 42.7 ...
```

We fit and validate the additive model:

```
m1 <- lm(volume ~ lot + trough, data = sphagnum)
```

```
par(mfrow = c(2, 2))
plot(m1)
```



```
par(mfrow = c(1, 1))
```

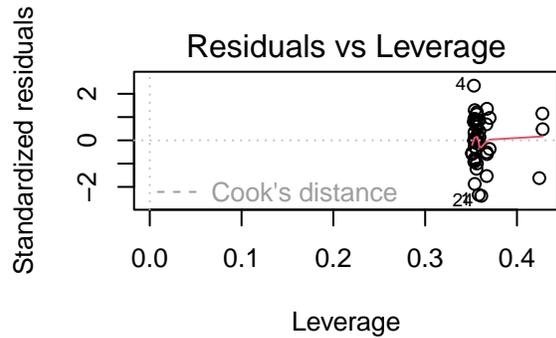
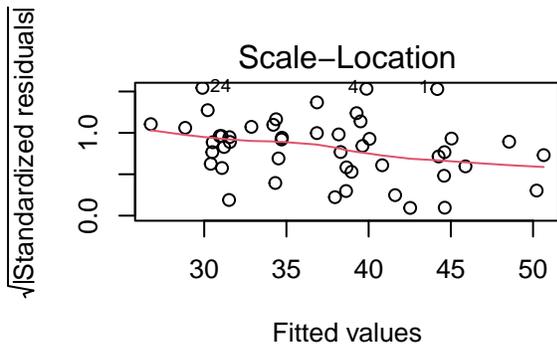
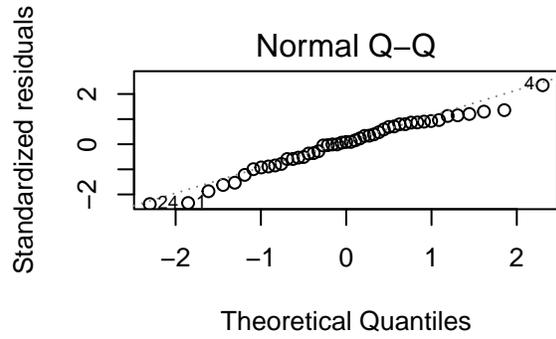
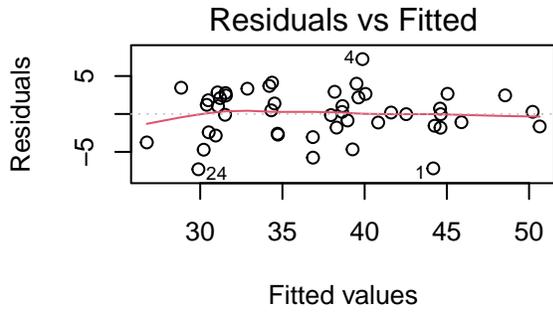
We see that observation 46 is an outlier and investigate it:

```
sphagnum[46, ]
```

```
## lot trough volume
## 46 12      2    7.5
```

We remove this outlier (be careful about this in practice!) and refit the model:

```
m2 <- lm(volume ~ lot + trough, data = sphagnum[-46, ])
par(mfrow = c(2, 2))
plot(m2)
```



```
par(mfrow = c(1, 1))
```

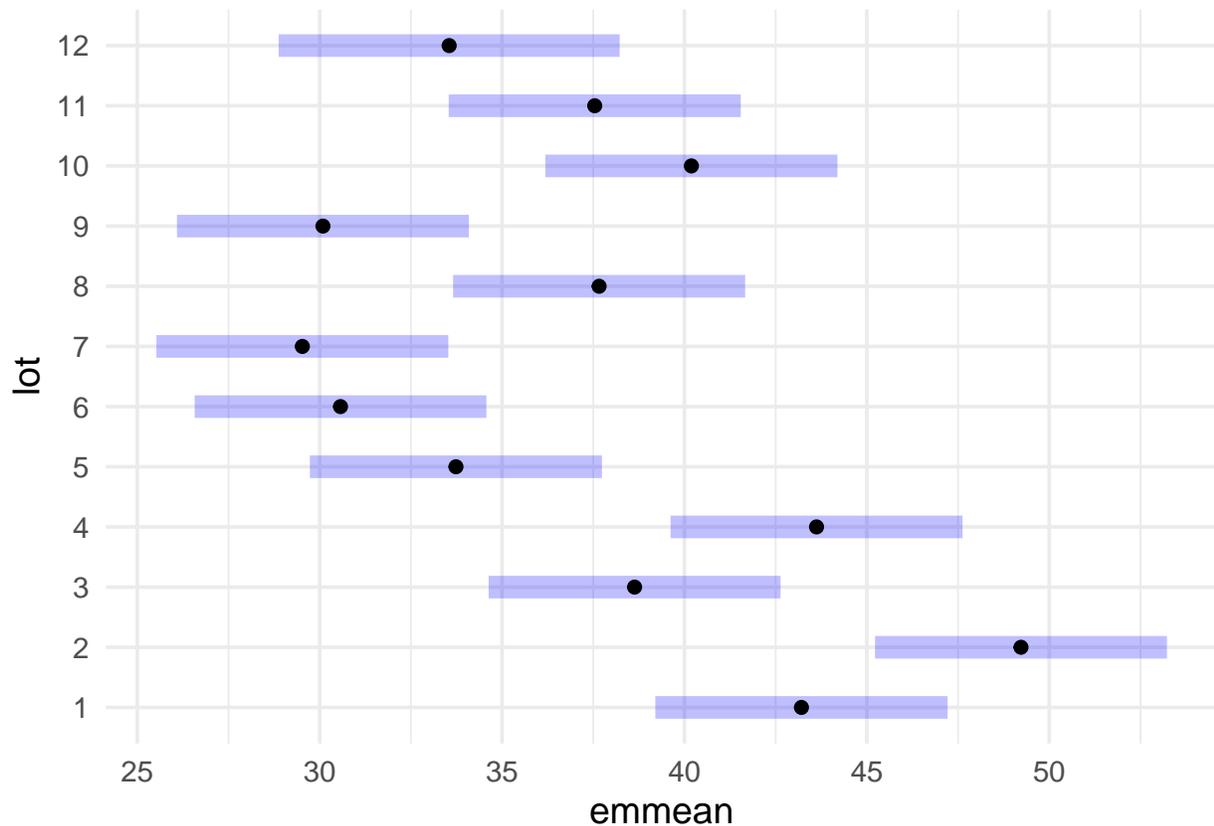
The plots all look okay! We now use obtain the desired em-means:

```
emmeans(m2, ~lot)
```

```
##   lot emmean   SE df lower.CL upper.CL
## 1    43.2 1.96 30    39.2    47.2
## 2    49.2 1.96 30    45.2    53.2
## 3    38.6 1.96 30    34.6    42.6
## 4    43.6 1.96 30    39.6    47.6
## 5    33.7 1.96 30    29.7    37.7
## 6    30.6 1.96 30    26.6    34.6
## 7    29.5 1.96 30    25.5    33.5
## 8    37.7 1.96 30    33.7    41.7
## 9    30.1 1.96 30    26.1    34.1
## 10   40.2 1.96 30    36.2    44.2
## 11   37.5 1.96 30    33.5    41.5
## 12   33.5 2.29 30    28.9    38.2
##
```

```
## Results are averaged over the levels of: trough
## Confidence level used: 0.95
```

```
plot(emmeans(m2, ~lot))
```



We see that the em-means obtained here are different from the ones computed in the given table. This is because the `emmeans` function chooses a balanced value for the remaining model covariates when comparing the means in each lot. The em-means therefore give a better idea of the mean in each lot that is comparable, since we are controlling for the effect of the watering trough.

## Exercise 4.4 Linear regression

### Problem

In a field experiment the concentration of phosphorus available for plant growth was measured for each of 18 plants. Furthermore, the concentration of inorganic phosphorus was chemically determined and the concentration of an organic phosphorus component was measured for each plant. The primary interest of the study is to describe the concentration of phosphorus available as a function of the concentrations of inorganic and organic phosphorus. We have the following *Table of Variables*:

Variable	Type	Usage
inorganic	continuous	fixed effect
organic	continuous	fixed effect
available	continuous	response

The dataset is shown below, and it is also available in the text file `phosphorus.txt`:

inorganic	organic	available
0.4	53	64
0.4	23	60
3.1	19	71
0.6	34	61
4.7	24	54
1.7	65	77
9.4	44	81
10.1	31	93
11.6	29	93
12.6	58	51
10.9	37	76
23.1	46	96
23.1	50	77
21.6	44	93
23.1	56	95
1.9	36	54
26.8	58	168
29.9	51	99

Analyze the data, i.e. answer the generic questions:

- Is there an association?
- What is the association?
- Can the conclusions be trusted?

Hints and suggestions: If you do a multilinear regression of **available** on **inorganic** and **organic**, then one of the observations is not well-modelled. You may either decide to remove this observation (what is the easiest way to do this in R?). Alternatively, you may try a different analysis e.g. by doing a logarithmic transformation of the response variable.

(Reference: Exercise 8.4 from Anders Tolver & Helle Sørensen: *Lecture notes for Applied Statistics*.)

### Solution

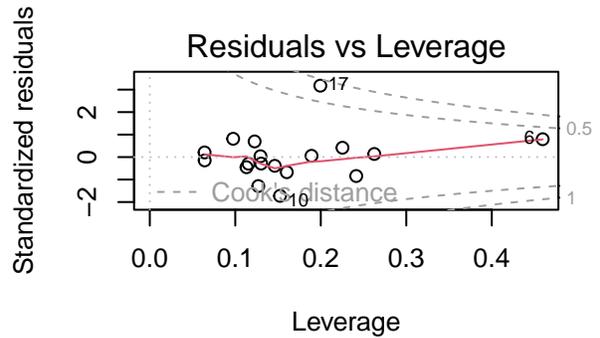
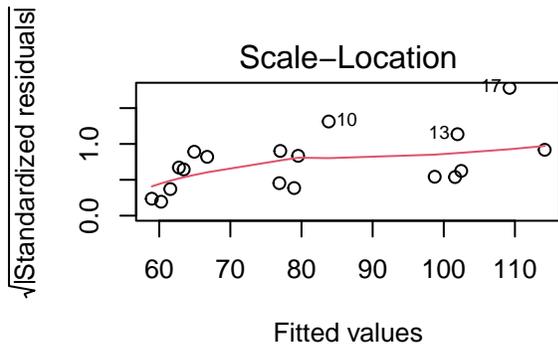
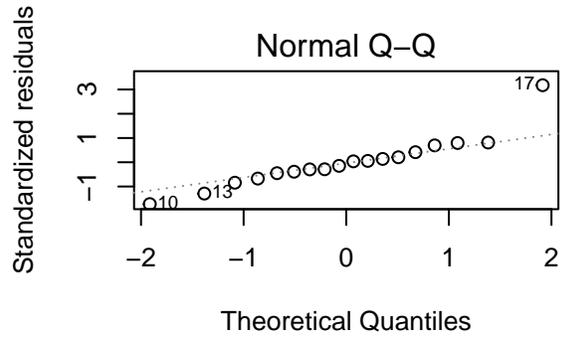
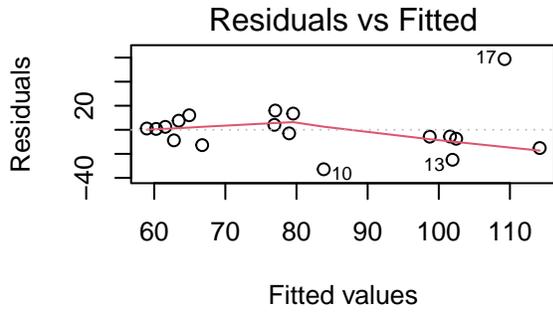
We read the data and summarize it:

```
phosphor <- read.delim("phosphorus.txt")
summary(phosphor)

##      inorganic      organic      available
## Min.   : 0.40   Min.   :19.00   Min.   : 51.00
## 1st Qu.: 2.20   1st Qu.:31.75   1st Qu.: 61.75
## Median :10.50   Median :44.00   Median : 77.00
## Mean   :11.94   Mean   :42.11   Mean   : 81.28
## 3rd Qu.:22.73   3rd Qu.:52.50   3rd Qu.: 93.00
## Max.   :29.90   Max.   :65.00   Max.   :168.00
```

We fit a multilinear regression and validate it:

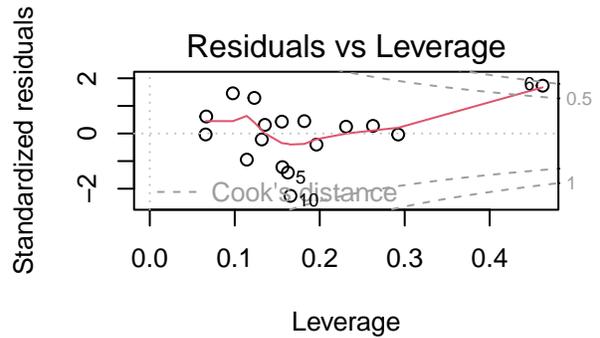
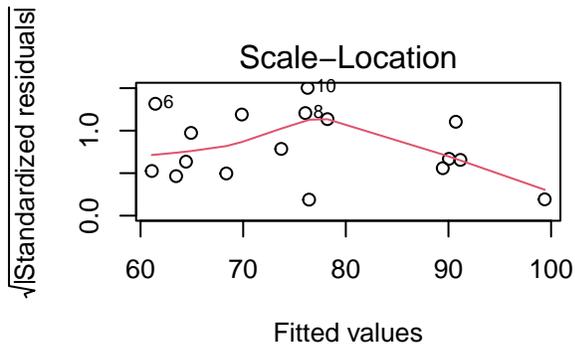
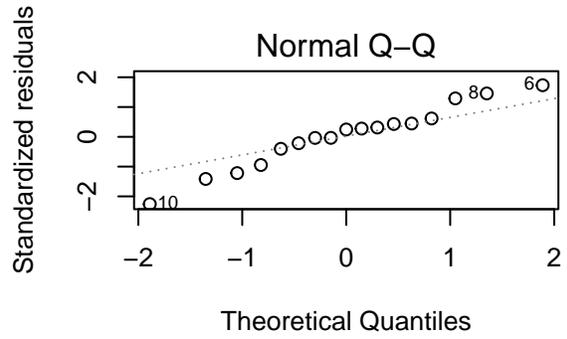
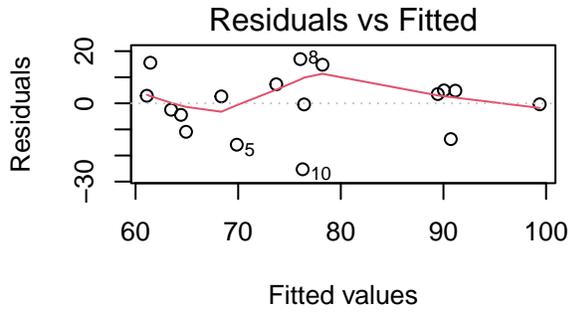
```
m1 <- lm(available ~ inorganic + organic, data = phosphor)
par(mfrow = c(2, 2))
plot(m1)
```



```
par(mfrow = c(1, 1))
```

Observation 17 is an outlier. We can fit the model without this observation:

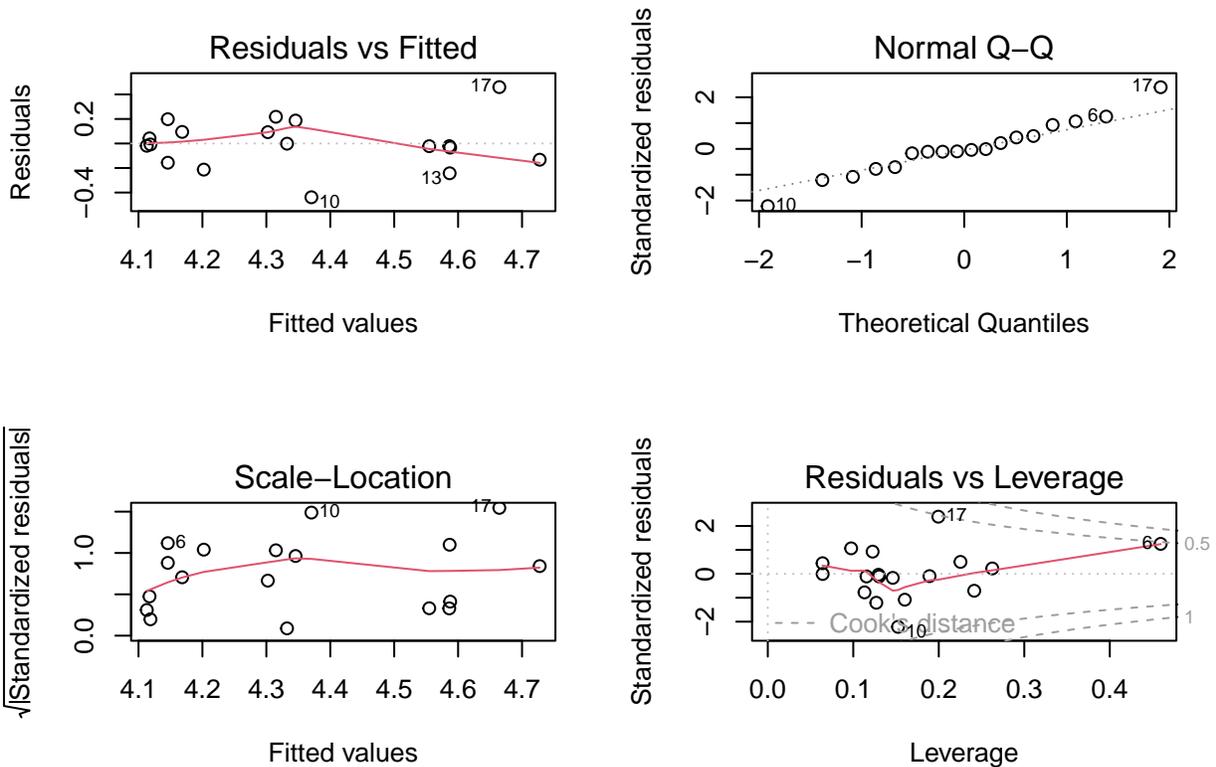
```
m2 <- lm(available ~ inorganic + organic, data = phosphor[-17, ])
par(mfrow = c(2, 2))
plot(m2)
```



```
par(mfrow = c(1, 1))
```

This looks much better! However, it is bad practice to remove an outlier without a good reason to do so, so we instead try a log transformation:

```
m3 <- lm(log(available) ~ inorganic + organic, data = phosphor)
par(mfrow = c(2, 2))
plot(m3)
```



```
par(mfrow = c(1, 1))
```

This has also mitigated the problem and the model looks acceptable. We test whether there is an effect of the covariates:

```
drop1(m3, test = "F")
```

```
## Single term deletions
```

```
##
```

```
## Model:
```

```
## log(available) ~ inorganic + organic
```

```
##           Df Sum of Sq    RSS   AIC F value    Pr(>F)
```

```
## <none>                 0.69337 -52.618
```

```
## inorganic  1  0.59109  1.28446 -43.521 12.7875 0.002758 **
```

```
## organic    1  0.00005  0.69342 -54.617  0.0011 0.974301
```

```
## ---
```

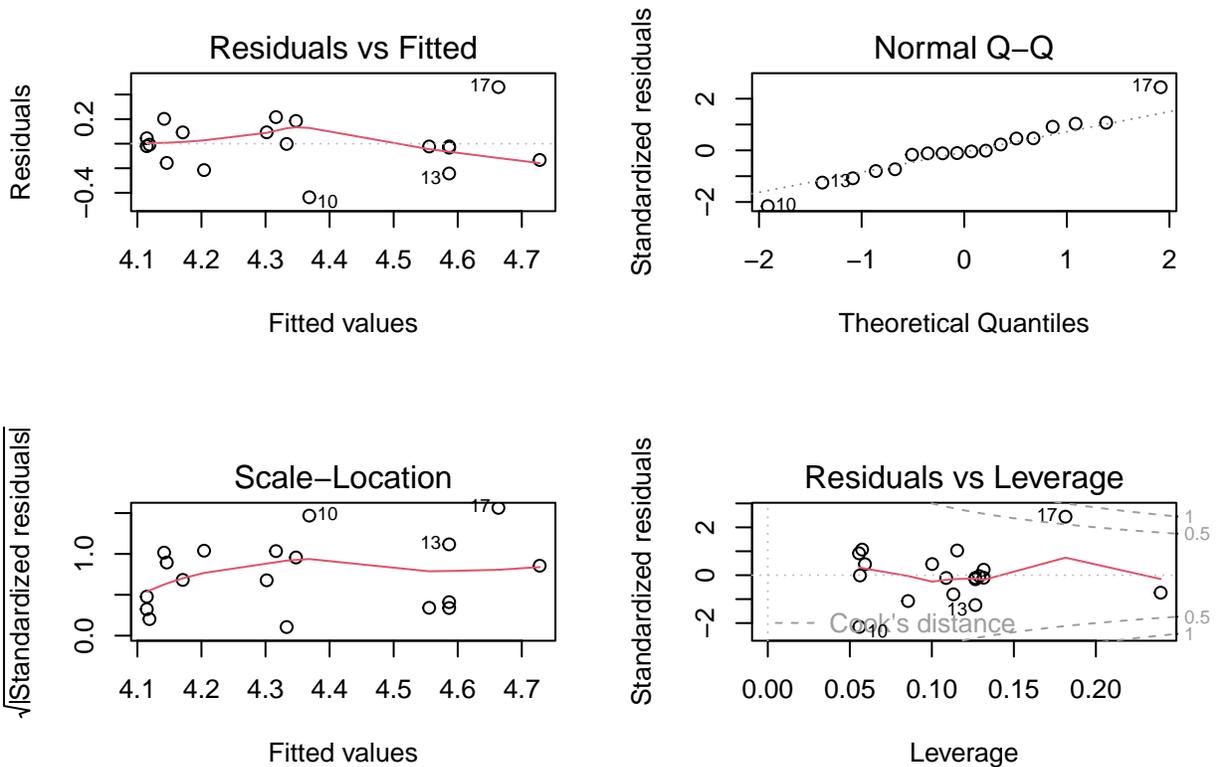
```
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

It seems that `organic` is not important; we refit the model without this variable, validate it and check for an effect of `inorganic`:

```
m4 <- lm(log(available) ~ inorganic, data = phosphor)
```

```
par(mfrow = c(2, 2))
```

```
plot(m4)
```



```

par(mfrow = c(1, 1))
drop1(m4, test = "F")

## Single term deletions
##
## Model:
## log(available) ~ inorganic
##           Df Sum of Sq    RSS   AIC F value    Pr(>F)
## <none>          0.69342 -54.617
## inorganic  1    0.75748  1.45090 -43.327  17.478 0.0007065 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

```

The model still looks okay and we see that `inorganic` is highly significant with a  $p$ -value of around 0.0007. We can construct a confidence interval for coefficients:

```

cbind(estimate = coef(m4), confint(m4))

##           estimate    2.5 %    97.5 %
## (Intercept) 4.10639987 3.9430857 4.26971405
## inorganic   0.02078734 0.0102467 0.03132798

```

Since we are modelling on a logarithmic scale, it is useful to backtransform the coefficient for `inorganic`:

```

exp(c(exp = coef(m4)[2], confint(m4)[2, ]))

## exp.inorganic    2.5 %    97.5 %
##           1.021005    1.010299    1.031824

```

An increase of 1 inorganic phosphorus will increase the available phosphorus by 2.1%.